



US 20030033324A1

(19) **United States**

(12) **Patent Application Publication**  
**Golding**

(10) **Pub. No.: US 2003/0033324 A1**

(43) **Pub. Date: Feb. 13, 2003**

(54) **RETURNING DATABASES AS SEARCH RESULTS**

(22) Filed: **Aug. 9, 2001**

(76) Inventor: **Andrew R. Golding**, Waltham, MA (US)

**Publication Classification**

(51) **Int. Cl.<sup>7</sup> ..... G06F 7/00**


(52) **U.S. Cl. .... 707/200**

Correspondence Address:  
**FISH & RICHARDSON PC**  
**225 FRANKLIN ST**  
**BOSTON, MA 02110 (US)**

(57) **ABSTRACT**

A method including enumerating plausible queries of a target database using query generation rules, and generating associated teasers for each of the enumerated queries using query matching rules.

(21) Appl. No.: **09/927,015**






FIG. 1




FIG. 2

100




FIG. 3

## RETURNING DATABASES AS SEARCH RESULTS

### TECHNICAL FIELD

[0001] This invention relates to returning databases as search results.

### BACKGROUND

[0002] A Web-enabled search engine is software that receives search terms from a user and identifies documents on the Web that contain, or are otherwise associated with, those search terms. Conventional search engines create an inverted index of all the terms that occur in all the documents that have been spidered on the World Wide Web. The search engine receives a user query and attempts to match the user query with terms in the index. Uniform resource locators (“URLs”) represent addresses of Web locations that contain hyperlinks that contain the full text of the identified documents within the index.

### SUMMARY

[0003] In general, in one aspect, the invention features a method including enumerating plausible queries of a target database using query generation rules, and generating associated teasers for each of the enumerated queries using query-matching rules.

[0004] Implementations of this aspect may include one or more of the following features.

[0005] The method may include storing the enumerated queries and their associated teasers in a lookup table. The method may also include receiving a user query of the target database, determining whether the user query matches an enumerated query in the lookup table, and displaying the teaser associated with the enumerated query in response to determining. The query generation rules may be domain specific. The query matching rules may be domain specific. Generating may also include conflict resolution rules. The target database may reside on a server connected to the Internet.

[0006] In general, in another aspect, the invention features a method including identifying queries that match elements in a target database, receiving a user query, determining if the user query matches one of the identified queries, and if the user query matches one of the identified queries, providing target database information to a user that relates to the user query.

[0007] Implementations of this aspect may include one or more of the following features.

[0008] The database may reside on a server. The server may reside in a network. Identifying may include applying query-generation rules to the target database, applying query-matching rules to each of the queries to generate associated teasers, and building a mapping from the queries to their associated teasers. Building a mapping may include storing the queries and associated teasers in a hash table. Building a mapping may include storing the queries and associated teasers in a cache. Building a mapping may include storing the queries and associated teasers in a trie data structure.

[0009] In general, in another aspect, the invention features a method including pre-processing a target database, build-

ing a mapping from selected queries to associated teasers for the target database, receiving a user query for the target database, and returning an associated teaser if the user query matches one of the selected queries.

[0010] Implementations of this aspect may include one or more of the following features.

[0011] Pre-processing may include identifying selected queries in conjunction with query-generation rules, and generating an associated teaser for each of the selected queries in conjunction with query-matching rules. Building a mapping may include storing each of the selected queries with the associated teaser. Storing may include placing each of the selected queries with associated teaser in a trie data structure. Storing may include placing each of the selected queries with associated teaser in a hash table. Storing may include placing each of the selected queries with associated teaser in a cache. Storing may include placing each of the selected queries with associated teaser in a lookup table. The method may also include displaying the associated teaser.

[0012] Embodiments of the invention may have one or more of the following advantages.

[0013] The method enables receipt of an arbitrary user query input, decides whether it matches the database, and if it does, returns a description of the best matching parts of the database, displayed as teasers.

[0014] Because matching of the database can be slow if it is done intelligently using domain knowledge, the method may be split into two stages, i.e., a pre-processing stage in which a mapping from queries to teasers is built for all queries that match the database, and a run-time stage in which the query is looked up in the mapping.

[0015] The method entails matching the user’s query against an element in the database, and summarizing the best match found in a descriptive hyperlinked text string, or teaser. To minimize run-time computation, the method uses query-generation rules to anticipate all plausible queries that can match elements in the database. To match the query against database elements intelligently, the method uses query-matching rules. The query-generation and query-matching rules encode heuristic knowledge of the domain of the database.

[0016] The method not only returns search results per se, but also can be tied to other related tasks, such as targeted advertising of the contents of the database.

[0017] Other features and advantages of the invention will become apparent from the following description and the claims.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0018] FIG. 1 is a block diagram of a computer network.

[0019] FIG. 2 is a diagrammatic view of a process for returning a database as a search result in response to a user query.

[0020] FIG. 3 is a flowchart of a process for returning a database as a search result.

### DETAILED DESCRIPTION

[0021] Referring to FIG. 1, a computer network 10 includes a computer 12, such as a personal computer (PC),

connected to a communication network 14, such as the Internet, which uses TCP/IP (Transmission Control Protocol/Internet Protocol) or another suitable protocol. Connections in the network 10 may use Ethernet, a wireless link, a telephone line, and so forth. Communications network 14 contains a server 16, which may be a mainframe computer, a PC, or any other type of processing device.

[0022] The computer 12 contains a processor 18 and a memory 20. The memory 20 stores an operating system 24 such as Windows98® or Linux, a TCP/IP protocol stack 26 for communicating over the network 14, and a Web browser 28 such as Internet Explorer® or Netscape Navigator®, for accessing Web sites and pages hosted by devices such as server 16 on network 10.

[0023] The server 16 contains a processor 30 and a memory 32. The memory 32 stores machine-executable instructions 36, an operating system 38, a TCP/IP protocol stack 40 and a database 42. Instructions 36 may be part of an Internet search engine or not, and are executed by processor 30 to perform processes 50 (FIG. 2) and 100 (FIG. 3). That is, a user at computer 12 uses Web browser 28 to access server 16, which, in response, executes instructions 36 to perform the processes described below.

[0024] Referring to FIG. 2, there is shown a process 50 that returns the database 42 as a teaser in response to a user query. The database may be stored locally on server 16 or on any other server or combination of servers connected to the network 14. Often information contained on the World Wide Web (“the Web”) is structured not as a plain text document, but rather as a database. For example, there are restaurant databases and car databases and databases for almost anything. Searching databases involves different technologies than are used for searching text documents, since databases have greater structure and associated semantics than text documents. For example, when searching a database, one may use the knowledge that a particular field of the database contains a person’s name, and therefore use special name-matching heuristics to match the contents of the field to a user query. For instance, one might allow the query to match just the surname in the name field, or to match the surname followed by a comma followed by the given name. These kinds of matches cannot be done reliably in an unstructured text document. Process 50 includes a pre-runtime process 52 and a runtime process 54.

[0025] Pre-runtime process 52 enumerates all potential user queries that are reasonable queries to ask about the contents of the database 42 so that the system is ready for queries received in real time by the run time process 54. For each query, a “best” teaser, or short excerpt, to give of the database to describe the best results of the database for that query is produced. Databases contain content having multiple elements 44. Each of the elements represents content. Content may be provided in response to a user query. Enumerating all potential user queries is implemented in a query-generation process 56. The query-generation process 56 utilizes query-generation rules or heuristics to enumerate a set of queries that can plausibly be asked about each element in the database. The query-generation rules are stored in the memory of the server 16.

[0026] For example, if someone is trying to ask for a particular restaurant, the query generation rules suggest what phrases the user might use. Thus, for a seafood

restaurant, a user might give a generic phrase such as “seafood restaurant” or “fish”. Or the user might give a more specific query involving a name of a restaurant, such as “Legal Seafoods” or just “Legal”.

[0027] These query-generation rules are domain-specific, e.g., one needs to know something about restaurants to be able to generate restaurant queries. It is preferred that the query generation rules be derived using a learning procedure. The learning procedure works with queries that users have actually given, together with, for example, the restaurants that they have received as acceptable responses to these queries. Thus, one can record user sessions and notice that, for instance, if users type “seafood” as a query, the users end up clicking on “Legal Seafoods” as their notice of an acceptable response. That gives one a training example that lets one learn that one query that users type when they are interested in the restaurant “Legal Seafoods” is the query “seafood”. If the learning procedure is also given the information that “seafood” is the type of cuisine served by “Legal Seafoods,” then the learning procedure can use the above training example as evidence to infer the following query-generation rule: users may ask for a restaurant by entering its type of cuisine.

[0028] Documents tend to be open-ended pieces of information, whereas a database is a much more structured type of collection of information. Using knowledge about the domain of the database, one can determine the types of queries users may ask. In summary, for a Boston restaurant database example, the query-generation rules generate the name of a restaurant (and sub-strings of the name), its location, the type of cuisine, and perhaps the name of the head chef and his or her signature dish(es). The query-generation process 56 captures the types of phrases that users would tend to provide as queries if they were interested in the restaurant. The rules annotate each query they produce with the database element, and the field within that element, that the query was generated from. For example, the rules for a restaurant domain may specify that the query “legal” is generated from the restaurant name of the “Legal Seafoods” database element, the query “kendall square” is generated from the location of this element, and so forth.

[0029] The query-generation rules for a domain are developed by learning from training examples of users’ queries together with the database elements that satisfied the queries. Such training examples may be obtained in E-commerce applications, where users demonstrate that they are satisfied with a database element by purchasing it.

[0030] A teaser generation process 58 uses output from the query-generation process 56. Once the query generation process 56 has generated a query, the teaser generation process 58 finds the best excerpts from the database in response to that query. Here again domain specific knowledge is utilized. For example, in the restaurant domain, knowledge about the subject of restaurants can be used to do a more intelligent form of matching. If a user types the word “fish,” one can use knowledge of the restaurant domain to determine that the majority of restaurant names that match that query have the same cuisine, which in this example is seafood. One can therefore infer that the user is interested in seafood restaurants. This is an example of a domain specific inference.

[0031] The teaser generation process 58 utilizes query matching rules or heuristics. It is preferred to find all of the

database entries that match a query for any reason, or, in other words, they match by any of the different rules. Thus, the teaser generation process 58 generates the best teaser for each query enumerated by the query-generation process 56. More specifically, teaser generation process 58 takes an arbitrary user query as input, decides whether it matches the target database, and if it does, returns a description of the best-matching parts of the database. These descriptions are called teasers. For example, suppose a user enters the query “tomatoes”. An example of a teaser that might be returned to the user for a book database would be “Find How to Make Pasta Sauces and other books on Tomatoes.”

[0032] To find the best teaser for a query, the teaser generation process 58 collects all instances of that query, together with their annotations, that were produced by the query-generation process 56. For example, for the query “fish,” there might have been six instances generated:

- [0033] 1. “fish”: From name of Naked Fish Restaurant
- [0034] 2. “fish”: From name of Jimbo’s Fish Shanty
- [0035] 3. “fish”: From name of The Village Fish
- [0036] 4. “fish”: From name of Vaughan’s Fish & Chips
- [0037] 5. “fish”: From chef Michael Fish of Mike’s Barbecue
- [0038] 6. “fish”: From chef Michael Fish of Mike’s Barbecue II

[0039] Finding the best teaser for the query amounts to finding the most likely database element(s), if any, suggested by the data above, that a user who types the query “fish” would be looking for. This is accomplished in the teaser generation process 58 by applying query-matching rules. Query-matching rules include two parts: the rules themselves, which propose ways of matching the query against different fields of the database elements; and a conflict-resolution process, which controls how to break ties among the rules.

[0040] As an example of a query-matching rule for the Boston restaurant domain, consider the following:

- [0041] Dominant Cuisine Rule:
- [0042] IF most or all restaurants whose names match the query have the same cuisine (C)
- [0043] THEN propose teaser: “(R) and other Boston (C) restaurants”
- [0044] where (R) is the name of the most popular Boston restaurant with cuisine (C)

[0045] The “IF” portion of a rule is referred to as the antecedent portion of the rule. Applied to the six instances of the “fish” query above, the Dominant Cuisine Rule would notice that four out of four restaurants whose names match the query have seafood as their cuisine, and that therefore a plausible interpretation of the “fish” query is that the user is interested in Boston seafood restaurants. The rule further infers that the user will be interested in the most popular Boston seafood restaurant, which is taken here to be Legal Seafoods, even though this restaurant does not contain “fish”

in its name. The rule therefore suggests the teaser: “Legal Seafoods and other Boston seafood restaurants.”

[0046] After all query-matching rules have executed, a conflict resolution process 59 is invoked if necessary to break ties. The conflict resolution process 59 uses conflict resolution rules or heuristics to decide which of a number of query-matching rules provides the best interpretation of the user’s query. It is preferred that the conflict resolution rules involve a notion of popularity of the database elements that are being recommended. One simply compares the mass of evidence, or the total popularity of the evidence that is being suggested by each query-matching rule, and assumes that, a priori, the rule with the most evidence will be the one that will most likely satisfy the user. Thus, a natural strategy is to select the query-matching rule whose matching database elements, i.e., the elements that match the antecedent of the rule, have the greatest popularity. This selects the rule whose elements have, a priori, the maximum likelihood of being of interest to the user. In practice, popularity of database elements can be estimated from sales counts in the case of E-commerce domains, sales ranks assuming sales counts follow a known distribution, or click counts in the case of selectable items.

[0047] Concluding the “fish” example above, suppose that a second rule matches the six database elements above:

[0048] Matching Chef Rule:

- [0049] IF query matches the surname of a chef (C) at more than one restaurant
- [0050] THEN propose teaser: “(R) and (C)’s other Boston restaurants.”
- [0051] where (R) is chef (C)’s most popular restaurant

[0052] This rule suggests the teaser: “Mike’s Barbecue and Michael Fish’s other Boston Restaurants.” At this point, conflict-resolution process 59 is invoked to decide between the Dominant Cuisine Rule and the Matching Chef Rule. Using the maximum likelihood strategy mentioned above, presumably the total popularity of all seafood restaurants with “fish” in their name will exceed the popularity of Michael Fish’s restaurants. Thus the final teaser will be: “Legal Seafoods and other Boston seafood restaurants.”

[0053] A map building process 60 constructs a mapping from queries to teasers to display for each query. If no query-matching rule matches for a query, then the query is omitted from the mapping. The actual mapping may be represented internally as any data structure that supports fast retrieval, such as tries, hash tables or lookup tables. Caching results of lookup can further speed up retrieval.

[0054] During the run-time process 54, a user query is received from a user on the computer system 12 by a lookup process 62. The lookup process 62 determines whether the user query has been found in the mapping data structure. If the query is found, the teaser for that query, as given by the mapping, is sent to the user as a search result. The search result represents the database as an associated teaser that is displayed to the user. If the lookup process cannot find the user’s query in the mapping, no teaser is returned for display to the user.

[0055] Now referring to FIG. 3, there is shown a process 100 for returning a teaser representing a database as a search

result. A user query is received **102** by a lookup process. The lookup process determines **104** whether the user query is found in a mapping data structure. The mapping data structure represents a set of pre-processed user queries and associated teasers. If the user query is found in the mapping data structure, a search result containing an associated teaser that represents the database is returned **106** to the user for display. If the user query is not found in the mapping data structure, no teaser is returned **108** to the user in response to the query.

[**0056**] Other embodiments are within the scope of the following claims. For example, the method may also include presenting the teaser in conjunction with advertising. This advertising can be specifically targeted to the user by relating it to the user query. The query-matching rules can include additional text that relates to the proposed teaser.

What is claimed is:

1. A method comprising:
  - enumerating plausible queries of a target database using query generation rules; and
  - generating associated teasers for each of the enumerated queries using query-matching rules.
2. The method of claim 1 further comprising storing the enumerated queries and their associated teasers in a lookup table.
3. The method of claim 2 further comprising:
  - receiving a user query of the target database;
  - determining whether the user query matches an enumerated query in the lookup table; and
  - displaying the teaser associated with the enumerated query in response to determining.
4. The method of claim 1 wherein the query generation rules are domain specific.
5. The method of claim 1 wherein the query matching rules are domain specific.
6. The method of claim 1 wherein generating further comprises conflict resolution rules.
7. The method of claim 1 wherein the target database resides on a server connected to the Internet.
8. A computer program stored on a computer-readable medium, the computer program comprising instructions that cause a computer to:
  - enumerate plausible queries of a target database using query generation rules; and
  - generate associated teasers for each of the enumerated queries using query matching rules.
9. A method comprising:
  - identifying queries that match elements in a target database;
  - receiving a user query;
  - determining if the user query matches one of the identified queries; and
  - if the user query matches one of the identified queries, providing target database information to a user that relates to the user query.
10. The method of claim 9 wherein the database resides on a server.

11. The method of claim 10 wherein the server resides in a network.

12. The method of claim 9 wherein identifying comprises:
 

- applying query-generation rules to the target database;

applying query-matching rules to each of the queries to generate associated teasers; and

building a mapping from the queries to their associated teasers.

13. The method of claim 12 wherein building the mapping comprises storing the queries and associated teasers in a hash table.

14. The method of claim 12 wherein building the mapping comprises storing the queries and associated teasers in a cache.

15. The method of claim 12 wherein building the mapping comprises storing the queries and associated teasers in a trie data structure.

16. A method comprises:

pre-processing a target database;

building a mapping from selected queries to associated teasers for the target database;

receiving a user query for the target database; and

returning an associated teaser if the user query matches one of the selected queries.

17. The method of claim 16 wherein pre-processing comprises:

identifying selected queries in conjunction with query-generation rules; and

generating an associated teaser for each of the selected queries in conjunction with query-matching rules.

18. The method of claim 16 wherein building a mapping comprises storing each of the selected queries with the associated teaser.

19. The method of claim 18 wherein storing comprises placing each of the selected queries with associated teaser in a trie data structure.

20. The method of claim 18 wherein storing comprises placing each of the selected queries with associated teaser in a hash table.

21. The method of claim 18 wherein storing comprises placing each of the selected queries with associated teaser in a cache.

22. The method of claim 18 wherein storing comprises placing each of the selected queries with associated teaser in a lookup table.

23. The method of claim 16 further comprising displaying the associated teaser.

24. A computer program stored on a computer-readable medium, the computer program comprising instructions that cause a computer to:

identify queries that match elements in a target database;

receive a user query;

determine if the user query matches one of the identified queries; and



provide target database information to a user that relates to the user query if the user query matches one of the identified queries.

**25.** The computer program of claim 24 wherein the instruction to identify comprises instructions to cause the computer to:

apply query-generation rules;

apply query-matching rules to each of the queries to generate associated teasers; and

build a mapping from the queries to the associated teasers.

**26.** A computer program stored on a computer-readable medium, the computer program comprising instructions that cause a computer to:

pre-process a target database;

build a mapping from selected queries to associated teasers for the target database;

receive a user query for the target database; and

return an associated teaser if the user query matches one of the selected queries.

**27.** An apparatus comprising:

a memory that stores executable instructions; and

a processor that executes the instructions to:

pre-process a target database;

build a mapping from selected queries to associated teasers for the target database;

receive a user query for the target database;

return an associated teaser if the user query matches one of the selected queries; and

display the associated teaser to the user.

**28.** An apparatus comprising:

a memory that stores executable instructions; and

a processor that executes the instructions to:

enumerate plausible queries of a target database using query generation rules; and

generate associated teasers for each of the enumerated queries using query matching rules.

\* \* \* \* \*